Selling Subscriptions*

Liran Einav, Ben Klopack, and Neale Mahoney†

August 3, 2023

Abstract

Retailers are increasingly selling goods and services via subscriptions instead of spot markets. In this paper, we study one benefit to the retailer of selling subscriptions: the possibility that – presumably because of inattention or inertia – consumers continue to pay for subscriptions after the flow benefit falls below its price. We use comprehensive data from a large payment card network and focus on credit and debit cards that get replaced (e.g., due to expiration). Replaced cards require an active subscription renewal decision, and we document that months during which cards are replaced are associated with much higher rates of cancellation for the ten subscriptions we study. We write down and estimate a stylized model of subscription renewals that allows us to recover the baseline degree of inattention. We find that estimated inattention is higher for consumers that took cash advances, a proxy for low financial sophistication. Relative to a counterfactual in which consumers are fully attentive, inattention raises seller revenues by between 14% and more than 200%. We use the estimated model to explore the quantitative impact of possible regulatory remedies.

*We are grateful to Jean-Pierre Dubé for helpful suggestions, and to seminar participants at Stanford, Tel Aviv, TxIO Conference, and Montreal IO conference for useful comments. We thank Chris Xue and especially Katja Hofmann for superb research assistance.

†Einav: Stanford University and NBER, leinav@stanford.edu; Klopack: Texas A&M, bklopack@tamu.edu; Mahoney: Stanford University and NBER, nmahoney@stanford.edu. Mahoney was on leave from Stanford and NBER at the White House in 2022-2023; all his work on this paper was done while he was at Stanford.

1 Introduction

A growing number of retail goods and services are now marketed as recurring subscriptions, which are typically billed at regular (often monthly) frequencies. The goods and services are then provided indefinitely until a consumer actively cancels. While subscriptions have been pervasive in some product categories for many years (e.g., newspapers or gym memberships), their use has expanded more recently to digital products (e.g., media streaming services or software licenses), home security systems, consumer products (clothing, shaving, and makeup needs), and ingredients for home-cooked meals. According to some estimates, the size of the "subscription economy" more than quadrupled over the last decade (Zuora 2022).

This rapid growth is often attributed to two primary factors. On the supply side, digital products have become a larger share of the retail sector, and such products may lend themselves more naturally towards a subscription model. On the demand side, there seems to have been an increased emphasis on convenience, and subscriptions are often associated with more convenient, hassle-free transactions.[1]

In this paper, we explore the potential for a third factor to play a quantitatively important role in the growth of the subscription economy. While subscription products may provide convenience benefits to consumers, they may also allow firms to exploit inattentive subscribers, and thus incentivize firms to shift their payment model from pay-as-you-go to subscription models. Because subscriptions are automatically renewed, consumers who are inertial or not fully attentive may continue to pay for services they no longer value. Indeed, there are now multiple new companies whose business model (marketed as a subscription!) is to help subscribers find and cancel unwanted subscription services.[2] If consumers do not fully anticipate their inertia at sign-up, this may create supply-side incentives that amplify the growth of subscription offerings.[3]

To quantify these incentives, we use transaction-level data from a large payments card company to analyze consumer renewal and cancellation behavior in the context of ten popular subscription services in the US. The key research design underlying our exercise takes advantage of our ability to observe card replacement (when cards expire, are lost, or stolen), and to link new cards to the cards they replaced. Because the replacement card is associ-

[1] Surveys report that 32% of US consumers "signed up to the subscription because it feels nice to receive something every month" (Emarsys 2021). A decade ago, a similar emphasis on increased convenience may have helped explain the transition on eBay from auctions to fixed prices (Einav et al. 2018).

[2] A recent survey found that nearly 90% of consumers underestimated their monthly spending on subscriptions, with the average respondent spending more than three times their initial estimate (West Monroe 2021).

[3] This motivation is also central to the recent FTC complaint against Amazon surrounding its Prime subscription service (Federal Trade Commission 2023a).

ated with a new card number, consumers typically need to update their billing information with the subscription provider, inducing an active renewal choice. Indeed, we document a sharp, abnormal drop in subscriber retention rates during the month of card replacement. This empirical pattern is present for all subscription services that we study. We also verify that overall transaction activity on the card remains almost the same, which rules out the possibility that the sharp drop in retention rates is driven by subscribers switching their subscriptions to a different card or by some other change in broader purchasing behavior.

In order to quantitatively interpret this pattern and assess the impact of inertia or inattention on seller revenues, we write down and estimate a stylized model of subscriber renewal behavior. In the model, subscribers are myopic and sign up for a service when their monthly flow utility from the service is greater than the monthly price. We then assume that the flow utility follows an AR(1) process, and that a fully attentive subscriber would cancel their subscription as soon as their flow utility falls below the monthly fee. Yet, we assume that in most periods (with the exception of the month of card replacement) subscribers are imperfectly attentive, and make an active choice with probability $\lambda < 1$, leading them to continue paying for the service even when they "shouldn't."

We estimate the model for each service separately, and despite only having three parameters (for each service), it replicates the key patterns in the data remarkably well. Consistent with the sharp drop in retention rates during the month of card replacement, we estimate a fairly large degree of inattention, with estimates of λ that vary (across services) between 0.04 and 0.5. We also show that consumer inattention is higher for consumers that took cash advances, who may be less financially sophisticated.

We then use the model to perform counterfactual exercises that assess how much faster (on average) subscribers would cancel their subscription if they were fully attentive ($\lambda = 1$). We find that seller revenues (or equivalently average subscription durations) are significantly higher due to subscriber inattention, with important heterogeneity across services. We estimate that inattention increases firm revenues by between 14% and more than 200%, depending on the service. We then use the estimated model to explore the potential impact of simple policy remedies, which – in the spirit of recent policy guidance from the Federal Trade Commission (2021) – would require all subscription services to "force" an active renewal decision at regular frequencies.

Our empirical analysis focuses on the intensive margin decision (renewal or cancellation), taking as given the set of subscribers who decided to initially subscribe. In order to fully understand the ability of service providers to exploit imperfectly attentive subscribers, one must also assess whether consumers anticipate their future inattention when they initially decide to subscribe. This question is studied in important complementary work by Miller,

Sahni, and Strulov-Shlain (2023), in the context of a newspaper subscription in Europe. Using a field experiment that varies whether a contract automatically renews or cancels, the authors show that consumers consider their future inattention: they are less likely to begin subscribing under auto-renew, suggesting that subscribers partially anticipate their inertia. Apart from our focus on cancellation, our work also differs from Miller, Sahni, and Strulov-Shlain (2023) in that we study multiple services and document significant heterogeneity both across services and across consumer types.

We are obviously not the first to document consumer inertia or inattention. One of the earliest papers in this literature is the seminal study of DellaVigna and Malmendier (2006), who analyze the contractual choice and attendance pattern of health club members and document substantial cancellation lags for consumers enrolled in contracts with automatic monthly renewal. Another example is a more recent study, in which Posner et al. (2022) measure the impact of auto-renewal in political contributions in the context of the 2020 elections. They show that making auto-renewal the default option substantially increased the amount of donations, but they also show that auto-renewal increased the rate of refund requests, suggesting that not all of the incremental donations were "intentional."[4]

Our paper is also closely related to the literature on optimal contract design in light of "behavioral" consumers, who may exhibit inattention or limited self-control. DellaVigna and Malmendier (2004) show that firms may exploit the overconfidence of present-biased consumers in making an active cancellation decision and design contracts with a back-loaded pricing structure and automatic contract renewal. Accounting for heterogeneous consumer types, Johnen (2019) studies how firms trade off the exploitation of naïvely inattentive consumers and the adverse selection of sophisticated consumers, who make an active decision about contract renewal and can avoid high renewal prices.

Lastly, our paper also relates to a growing literature in computer science and law that studies the prevalence and impact of deceptive user interfaces on websites and smartphone apps, sometimes referred to as dark patterns (Mathur et al. 2019; Di Geronimo et al. 2020). In experimental studies, Luguri and Strahilevitz (2021) show that hidden information, long obstruction conditions, and trick questions are effective strategies to prevent consumers from rejecting automatic sign-up to a subscription service. Our paper is also more broadly related to policy efforts to reduce inertia in subscription plans, such as the "click to cancel" rule proposed by the Federal Trade Commission (2023b).

The rest of the paper is organized as follows. In section 2 we describe the data, the selection of subscription services, and the construction of the sample. In section 3 we present

[4]Two other related papers – Reme, Røhr, and Sæthre (2022) and Ascarza, Iyengar, and Schleicher (2016) – study the dynamics of inattention and subscription attrition.

descriptive evidence that motivates our key exercise and illustrates our empirical strategy. Section 4 presents the model and its estimation. In section 5 we show the estimation results and the counterfactual exercises, which help us to quantitatively assess how inattention affects firm revenues. The final section concludes.

2 Data and sample construction

2.1 Data

Data source. Our primary data source covers transactions conducted on a large payment card network in the United States between August 2017 and December 2021. For the set of subscription services used in our baseline sample (described below) and using publicly available information on total subscribers, we estimate that our data covers approximately 30% of all subscribers.

An observation is a single transaction, and the information on each transaction is similar to the typical line one would find on monthly credit card statements: the name of the merchant, a unique card identifier, a transaction amount, and a date. Importantly, there is no information on the specific goods or services that were purchased nor their prices. The sample is depersonalized and does not contain the name, address, or any other personally identifiable information about the cardholder.

Critically for the empirical strategy we describe below, a card in the data is associated with a unique account identifier, so that multiple cards within the same account can be linked. Thus, if a card expires or is lost or stolen, its replacement card has a new card identifier but retains the same account identifier. Because the quality of the account identifier variable is low for cards that were replaced in the early part of our sample, we focus our analysis on cards that are replaced in July 2018 and after. We note that we cannot link multiple accounts held by the same consumer, so we treat accounts as independent of each other.

Identifying subscription services. In order to systematically focus on a set of subscription services, we consider an industry report as a starting point. Specifically, West Monroe (2021) surveyed respondents about their spending on subscriptions in 21 different categories with 49 examples of specific services. We augmented this list by searching for industry reports for each category and adding any additional services with more than 500,000 subscribers (as reported by public sources). This process yielded a list of 69 subscription products. Of this set, we were able to identify (by "manually" name searching) 57 of these services in the

payment card data.[5]

We imposed the following additional criteria to arrive at a final list of subscription products. First, we required that services have a minimum of 500,000 average monthly subscribers in our data, which eliminates 31 of the 57 services. We then dropped 4 services that are primarily sold in long-term contracts (two cell phone and two internet service providers), 6 services that are sold by merchants with many non-subscription products,[6] 2 services with average subscription length shorter than six months, 2 services that were launched toward the end of our observation period, and 2 services with non-monthly billing.

Our final sample thus includes a set of 10 large subscription services that we use for our analysis. We identify them throughout by letters (A through J) as our data use agreement prevents us from revealing the merchant names. The 10 services cover both digital and non-digital products from several different merchant categories, including entertainment, security, retail goods, and newspapers.

The product offerings of these 10 services remained largely stable during our sample period, with relatively small changes in pricing.[7] We investigated consumer response to these price changes in the data and found it to have a negligible impact on consumer cancellation rates, so we abstract from these (small) changes for the rest of the analysis. However, we use service-specific month fixed effects in our empirical specification, which absorbs in principle any remaining impact of changes in prices or product offering.

2.2 Sample construction

Our primary research design takes advantage of the behavior of consumers that have their payment card replaced within the same account. We thus limit our sample to consumer credit and debit accounts[8] that had their cards replaced exactly once between July 2018 and January 2021. There are about 23 million accounts (and about 35 million account-service pairs) that meet these criteria and transacted with at least one of the ten subscription services analyzed.

We organize the data at the monthly level, and use the last transaction made on the old card to denote the last month in which the old card was available, meaning that the

[5]We can provide additional details on the process that led to selecting these 69 services upon request, and subject to review by the data provider.

[6]Because we are unable to observe which items were included in a transaction, a merchant that sells both subscriptions and other products would make it hard for us to identify subscribers separately from other customers.

[7]Two of the ten services made one-time price increases of $1-$2 to the monthly rate of the base subscription package, a third service increased the monthly price of their family package by $1, and a fourth service reduced the price of their base package, while increasing the price of their premium services.

[8]This excludes card types intended for business use and prepaid cards.

subsequent month is defined as the first month in which the new card had become active. We drop accounts for which there was a gap of more than one month between the last time the old card was used and the first time the new card was used (7% of the sample of 23 million accounts mentioned above) and accounts where the old card continued to be used after the replacement card was issued (an additional 12% of accounts). These restrictions leave us with a set of about 19 million accounts and approximately 28 million account-service pairs.

To construct our final sample, we make further restrictions to ease the analysis and graphical exposition of our results. Specifically, we include an account-service pair in the final sample if the card replacement occurs exactly 6, 12, or 18 months[9] after the first month of subscription to the service.[10] Once included in the sample, we track each account-service pair for 25 months and require that the account is active – that is, that the account is associated with at least one transaction (any transaction) in each of the 25 months (accounts that do not satisfy this activity requirement are excluded from the analysis sample). The resulting analysis sample is a relatively small subset – 870,358 account-service pairs, representing 800,545 distinct accounts – of these 28 million account-service pairs.[11]

Our final sample thus contains a collection of cohorts of initial subscribers to each subscription service. A cohort (of subscribers to a particular service) is defined by an initial subscription month and a card expiration at month x, where x is equal to 6, 12, or 18 months. The initial subscription month runs from January 2018 through July 2019.[12] For a given service, we therefore observe a total of 57 cohorts: 19 cohorts defined by their initial subscription months, each partitioned to three sub-cohorts that are based on the number of months (6, 12, or 18) at which card replacement occurs.[13]

We apply two final "data cleaning" steps that facilitate the subsequent analysis. First, we guarantee that each account-service observation follows a simple data structure that would fit

[9]The choice of 6, 12, and 18 is our admittedly arbitrary attempt to evenly span the 25-month panel structure (see below) to facilitate graphical presentation of the data.

[10]Recall that the earliest month in our data is August 2017; to focus on initial subscriptions, we keep subscriptions that start in January 2018 and after.

[11]A subset of the accounts in our baseline sample are subscribed to more than one subscription service. While card replacement may lead to correlated cancellation decisions, we note that the data do not reveal such a pattern. That is, analyzing cancellation decisions in a sample of cardholders who are subscribed to two services does not suggest that the two cancellation decisions are more correlated during card replacement months relative to other months. We therefore treat each account-service pair as an independent observation throughout the paper.

[12]July 2019 is the latest month for which we can observe a card replacement that occurs 18 months after initial subscription.

[13]For four (out of the ten) services, we observe only 56 (rather than 57) cohorts because the relatively small service size and small number of card replacements we observe in 2018 lead to no observations associated with January 2018 subscribers whose card is replaced in July 2018.

a hazard model: If we observe an account transacting with a service but "skipping" a single month, we "fill in" that month,[14] and if we observe two months or more without transactions with the service, we assume that the cardholder unsubscribed to the service regardless of any subsequent transactions (which would likely represent "re-subscriptions").[15] Second, we exclude the first month we observe a transaction for each cohort. For most services, we observe a much larger drop in subscriptions after the first month than after subsequent months. We thus consider subscribers as more attentive immediately after the first month (relative to subsequent months) and view this first month as "special" – e.g., a "trial period" – and in what follows we consider the "initial" subscription month as the second month after we first see a transaction in our data. This restriction reduces the number of account-service pairs from 870,358 to 635,021.

3 Descriptive evidence

Empirical constructs. Consider a cohort (s, x), which is associated with a given service, an initial subscription in month s, and card replacement which occurs in month $s+x$. Denote the number of subscribers in each month $t \geq s$ by $N(t; s, x)$, and define the cohort-specific retention rate as

$$R_n(t; s, x) \equiv N(t; s, x)/N(s; s, x). \tag{1}$$

where $n \equiv t - s$ is the age of the cohort in months. That is, the retention rate is the share of initial subscribers that remain subscribed at age n.

The top panel of Figure 1 presents the data in its most granular form. It is focused on one subscription service ("service A") and only on the 19 cohorts whose cards are replaced 12 months after initial subscription ($x = 12$). For those 19 cohorts, we plot the retention rate, R_n, in each month, throughout the 24-month observation period. The pattern is quite similar across the cohorts, revealing a smooth decline in retention rates over time, with a sharp, abnormal drop in retention rates around the card replacement month (month 12).

Although the raw patterns across cohorts are quite similar, it seems natural to aggregate across cohorts in order to adjust for any possible differences in cohort sizes, service popularity, seasonal variation, and (relatively small, as mentioned earlier) changes in the product offering and monthly subscription prices.

[14] That is, if we observe no transaction in month $s + t$ to a given service, we still consider this account as "subscribed" as long as there are transactions in months $s + t - 1$ and $s + t + 1$. This adjustment is quantitatively small and raises the average subscription duration in the entire sample from 17.2 to 17.6 months.

[15] About 20% of accounts that unsubscribe for two months or more return to the service within the 25-month window.

To do so, we estimate the following regression (weighting by cohort size, $N(s;s,x)$) for each service separately:

$$R_n(t;s,x) = \beta_t + \gamma_{n,x} + \varepsilon_{t,s,x}, \tag{2}$$

where β_t is a month fixed effect and $\gamma_{n,x}$ is a fixed effect for the number of months since initial subscription, which is allowed to flexibly vary with x. We then define the adjusted retention rate by

$$\widehat{R}_n(x) \equiv \widehat{\gamma}_{n,x}/\widehat{\gamma}_{1,x}, \tag{3}$$

where the $\widehat{\gamma}_{n,x}$'s are the estimated coefficients from equation (2).

The bottom panel of Figure 1 illustrates this empirical construct – which we will deploy heavily below – in the context of the 19 cohorts shown in the top panel of the figure. It plots $\widehat{R}_n(x)$ for the same service ("service A") and x ($x = 12$), showing how the adjustment aggregates across cohorts and smooths out some of the cohort-specific noise.

Figure 1 is also the first time in which we see our inability to perfectly time the card replacement date. This is the reason why the sharp drop in retention rates around the month of card replacement does not happen in a single month, but instead spans two consecutive months. This pattern will repeat itself throughout, and we explain later how we adjust for it in the context of the estimation of the model of Section 4.

Appendix Table A1 summarizes the drop in retention rates during card replacement, reporting the average monthly change in retention rates ($\widehat{R}_n(x) - \widehat{R}_{n-1}(x)$) during the two-month replacement window and outside of it for each subscription service. On average, the monthly drop in retention is 0.08 during the replacement window, 4 times larger than the 0.02 drop during other months. This understates the difference since, as discussed above, we cannot pin down the exact month of card replacement.

Account activity around card replacement. Our economic interpretation (discussed in more detail below) of the sharp drop in retention rate around card replacement is that it reflects a "forced" active renewal decision by the cardholder, as opposed to a more passive, inertial renewal decision in other months. One key concern about this interpretation is that card replacement may have a broader impact on account activity, which means that the sharp drop in retention rate around card replacement is mechanical. For example, one plausible story is that the process of card replacement is associated with short interruptions in access to the card (e.g., while waiting for the new card to arrive in the mail), which may make the cardholder switch to a different card. This may mean that the sharp drop in retention rate could simply reflect substitution of the subscription to this other account rather than cancellation.

In order to assess this concern, Figure 2 uses the entire analysis sample and presents the variation in account activity around the month of card replacement. The top panel shows the number of monthly transactions associated with the account; that is, the number of monthly transactions on the old card prior to replacement and the number of monthly transactions on the new card after replacement. The bottom panel repeats the same exercise but uses total monthly spending associated with the account.[16]

Figure 2 shows some disruption in account activity during the month of card replacement, but an almost immediate recovery to pre-replacement activity levels. The average number of monthly transactions and monthly spending fall, respectively, from 44.2 transactions and \$2202 in the month before card replacement to 37.6 and \$1928 (an 18% and 14% decline) in the month of card replacement, but recovers to 42.1 and \$2114 two months later.[17] We note that the fact that account activity after replacement remains 4% lower than pre-replacement level could be driven, fully or partially, by subscription lapses, and the attention interpretation below would be one mechanism that could lead to this incomplete recovery.

Cancellation rates around card replacement. Figure 3 presents the adjusted retention rates ($\widehat{R}_n(x)$) in the sample. Each panel presents retention rates for one service (the service is indicated by the letter at the top right corner of the panel), and includes three separate lines that correspond to the number of months between initial subscription and card replacement ($x = 6, 12, 18$).

The retention patterns are quite heterogeneous across the ten different services, but the common theme across all of them is a sharper drop in retention rates around card replacement. The sharper drop is noticeable but relatively small in some of the services (services C, G, and J) but is much larger in some of the others (e.g., services A, B, D, and I). Across all services the magnitude of the drop at card replacement is quantitatively similar for subscribers whose card is replaced 6 months, 12 months, and 18 months after initial subscription.

As mentioned earlier, we interpret these remarkably similar qualitative patterns across subscription services as indicating a more active renewal decision by subscribers as a result of card replacement, relative to a more passive and inertial renewal decision in other months. Despite the similar qualitative patterns, the quantitative drops in retention rates at card replacement are quite different across services. To provide a comparable quantitative metric for these drops, we specify a model of subscriber behavior in the next section. The model

[16]The figures look very similar if we exclude transactions at the ten subscription services we study.

[17]The fact that it takes two months rather than one to recover to the original level is due to our inability to perfectly time the month of card replacement, as discussed earlier.

allows for heterogeneity across services along other dimensions (e.g., in the decline in retention rates over time in non-replacement months), and also provides a way to quantify the importance of subscriber attention and the revenue impact of counterfactual policies that may alter retention rates. We turn to this next.

4 Model and estimation

Our primary objective is to estimate the revenue impact of inattention associated with the automatic renewal of subscription services. Because the focus is on the revenue impact for sellers rather than the utility impact on consumers, the model is static and highly stylized and should be viewed as a positive (rather than normative) description of renewal behavior. The key feature of the model we require for our exercise is the ability to simulate subscription behavior under varying degrees of consumer attentiveness.

A model of subscription renewal behavior. Consider a specific subscription service, which is associated with a monthly subscription price p,[18] and a potential subscriber i, whose flow utility from the service during month t is denoted by u_{it}.

We assume that u_{it} follows a Markov process, such that $u_{it} \sim F(\cdot|u_{i,t-1})$, and that subscribers – once they are already subscribed – are not taking into account any dynamic considerations, so their renewal decision only relies on the comparison between the flow utility u_{it} and the price p. This latter assumption is consistent with consumers being myopic or alternatively with consumers being forward-looking but failing to anticipate their future inattention.

Given these assumptions, all new subscribers must have $u_{it} > p$ for the month in which they subscribe to the service for the first time, so we normalize $t = 0$ for this initial subscription month, and denote by $G(u_{i0}|u_{i0} > p)$ the cross-sectional distribution of u_{it} for new subscribers.

In a typical month t, a subscriber can be either attentive or inattentive. If inattentive, the subscriber automatically renews the subscription. If attentive, the subscriber renews if and only if $u_{it} > p$. Subscribers are attentive in a given month with probability $\lambda_{it} \in (0, 1]$. Importantly, in the first month after card replacement, we assume subscribers are attentive with probability one because they are asked to actively enter the details of their new card.

Parameterization. We define the net flow utility as $v_{it} \equiv u_{it} - p$, and assume that $F(\cdot)$ is

[18] The estimation below is carried out on a service-by-service basis, so throughout we omit service subscripts for expositional clarity.

an AR(1) process (without a constant),

$$v_{it} = \rho\, v_{i,t-1} + \varepsilon_{it}, \tag{4}$$

where ε_{it} follows a mean-zero normal distribution with a standard deviation that is normalized to one. We assume that the distribution of initial net utilities – $G(u_{i0}|u_{i0} > p)$ or equivalently $G(v_{i0}|v_{i0} > 0)$ – is given by an exponential distribution,

$$v_{i0} \sim Exp(\eta)\ , \tag{5}$$

which has a mean and standard deviation η. Finally, in the baseline specification, we assume that the attention probability λ is the same across people and over time (for a given service).[19]

Taken together, the model can be summarized by three service-specific estimable parameters: the persistence and trajectory of the flow utility from the service (as given by ρ), the extent to which new subscribers are close to the renewal margin (η), and the overall level of attention or inertia (λ).

Estimation. We estimate the model separately for each service, by matching simulated moments. Specifically, we focus on the moments we have already seen in Figure 3: the adjusted retention rates $\widehat{R}_n(x)$ for each service, which vary by month since initial subscription n and by the timing of card replacement since initial subscription (that is, x is 6, 12, or 18 months). To account for the fact (mentioned earlier) that we cannot perfectly time the month of card replacement, we omit the month of card replacement from the set of moments we try to match.[20] Overall, for each service, we thus have 66 moments: For each of the three values of x (6, 12, and 18 months after initial subscription), we have 23 monthly retention rates, and we use all of them except the month in which the card is replaced. In estimating the parameters, we weight each moment by the corresponding cohort size ($\sum_s N(s; s, x)$).

To construct model predictions for a given set of parameter values, we use the model to simulate survival probabilities as a function of the three model parameters ρ, λ, and η (see Appendix A for more details), and estimate the parameters by minimizing the quadratic distance between the simulated moments and their empirical counterparts. While the parameters are allowed to vary flexibly across services, we require them to be the same for a given service across the three values of x (6, 12, and 18 months after initial subscription).

[19] We explore alternative assumptions in Section 5.3.

[20] For example, if card replacement occurs at month 6, Figure 3 shows the sharp drop in retention rates occurring over month 6 and month 7. By omitting the month-6 retention rate from the set of moments we match in estimation, we are essentially allowing the card replacement to occur in either month 6 or month 7.

Comparative statics and identification. We use the exercise shown in Appendix Figure A1 to build intuition for the identification of the model and illustrate its comparative statics. The figure uses the data and the estimated parameters associated with subscription service A.

The dashed lines in all panels report the model predictions for retention rates when the card replacement is in month 6 (left panels), month 12 (middle panels), and month 18 (right panels). We then change one parameter at a time (holding the other two at their estimated values) in order to hit an (arbitrary) retention rate of 0.56 by month 24, and explore how this affects the retention rates in earlier months. The two solid lines in each panel of Appendix Figure A1 show this exercise for a pair of parameters to facilitate comparison.

Consider for example the bottom left panel. In order to hit a retention rate of 0.56 in month 24 by *only* changing λ, the value of λ (that is, the attention level of subscribers) must be lower, leading to a fairly flat decline in retention rates before and after card replacement, and to a sharp drop in the retention rate on the month of card replacement. In contrast, if we wanted to hit the same level of retention rate in month 24 by *only* changing η, the value of η would have to be much higher so that more initial subscribers renew their subscription. Yet, this change in the parameter value would still predict a steep decline in retention rates before and after card replacement, and make the drop in retention rate on the month of card replacement much smaller.

Similar contrasts are illustrated in the other panels of the figure, which may help provide intuition for the separate identification of the three parameters. A flatter slope in retention rates before and after card replacement would load on λ, a steeper one would load on ρ and η, and variation in how retention rates change before versus after card replacement helps distinguish between the latter two (as illustrated, for example, in the middle center panel).

5 Results

5.1 Model fit and parameter estimates

The parameter estimates are shown in Table 1. Figure 4 presents the model fit, plotting the predicted retention rates from the estimated model against their empirical counterparts, service by service. In general, the fit of the model is – in our assessment – quite good, especially when taking into account the stylized nature of the model and the fact that it only has three parameters (for each service).[21]

[21] We note again that the model, by design, does not try to fit the fact that in the data the drop in retention rates covers two months rather than one. As discussed before, this is an artifact of our inability to "perfectly" time the card replacement in the data, and a pattern that we intentionally do not aim to

The parameter estimates are reasonably intuitive. Consider first the inattention parameter λ. A natural benchmark is $\lambda = 1$, which is a case where subscribers are fully attentive every period. Yet, the estimates of λ range (across services) between 0.044 (service I) to 0.5 (service G). For service I, almost all subscribers in a given month are inattentive and renew their subscription in a passive way, suggesting that inertia plays a major role and contributes significantly to the seller's revenues. This can be loosely seen in the shape of the retention rates for service I: a very flat pattern before and after card replacement, and a very sharp decline at card replacement, during which more than 30% of subscribers are lost. For service G, the empirical pattern is quite different: a steady and fairly steep decline in retention rates over time, and only a small incremental decline at the month of card replacement. Yet, even for this case, our estimate of λ is well below 1, implying that (on average) subscribers make an active decision only every two months on average.

The majority of the estimates of the ρ parameter are very close to 1, suggesting that preferences for the services are (on average) stable after initial subscription, and approximately follow a random walk. For several of the services, estimates of ρ are well below 1, implying that the service may find it difficult to retain consumers for long durations. These are services which may benefit the most from inattentive consumers.

Finally, the η estimates reflect the extent to which new subscribers who sign up to the service are mostly "marginal" subscribers who are at risk of quickly unsubscribing if attentive (low η) or mostly "infra-marginal" subscribers who would require a sequence of negative preference shocks before they cancel (high η). The estimates are quite heterogeneous across services, with service B drawing almost entirely "marginal" subscribers ($\eta = 0.003$) while services E, G, and H are associated with relatively high η estimates that are greater than 2.

5.2 Quantifying the impact of inattention

The impact of inattention on revenues. We now use the model to quantify the impact of consumer inattention on the revenues of the subscription. To do so, we use the model and the estimated parameters described above, and simulate the renewal decisions of a cohort of initial subscribers over 10 years (that is, 120 months). Importantly, for this exercise we assume that subscribers face no card replacement throughout, so they are attentive each period with probability λ. We then repeat the same exercise but assume fully attentive subscribers (that is, $\lambda = 1$) and compare the results. Throughout we make the simplifying assumption that subscribers who decide to not renew are lost "forever." This is a strong assumption, but can also be motivated by the observation that getting old subscribers to

replicate with the model.

resubscribe to the service could be almost as costly as attracting new "fresh" subscribers. We discount revenues at a rate of 1% per month.[22]

The results are shown in Table 2. The first column of the table shows the share of consumers that are unaffected by inattention. This group has positive valuations in every month and subscribes until the final period whether or not they are attentive. The remaining "affected" group of consumers have some negative realizations of v_{it} at some point during the period. The second column of the table shows the average number of months that this group subscribes when they are attentive with probability λ every period, and the third column shows the average number of months they subscribe when they are fully attentive in every month. The last column reports the ratio of the revenues (measured by the number of subscriber-months) the subscription service obtains (over a horizon of 10 years) when subscribers are attentive with probability λ relative to the revenues it would obtain from the same set of initial subscribers if they were attentive with probability 1 every month.

Overall, we find that the benefits from attention (measured by the revenue ratio) are substantial, yet highly heterogeneous across services. Inattention (relative to "perfect" attention) modestly increases revenues for some services (e.g., 14% for service G), but *triples* revenues for others (service B).[23] In other words, the average subscription duration for service B would drop from its observed duration of over a year to about 4 months if subscribers were fully attentive. It is plausible to suspect that this subscription service would not be viable from a business perspective if not for its subscribers' inattention.

Exploring the impact of possible remedies. One possible regulatory remedy to subscriber inattention is to "force" current subscribers to make an active decision every month. This could be implemented, for example, by requiring services to make it a default for subscribers to not renew the service unless they indicate otherwise. Clearly, this is an extreme intervention, which would likely result in some consumer backlash; after all, some of the benefits of subscription services are associated with the convenience of *not* having to actively renew the service on a monthly basis.

As a way to explore intermediate solutions, and in the spirit of recent FTC guidance (Federal Trade Commission 2021), we consider the effect of possible policy interventions that reduce inattention by inducing subscribers to make an active choice at regular intervals. The agency's Enforcement Policy Statement governing negative option marketing, which includes subscription products with automatic renewal, free trials that are converted to paid subscriptions, and similar sales strategies, requires firms to ask for informed consent for a

[22]The revenue benefits from inattention are not very sensitive to discounting.

[23]Appendix Figure A2 provides a more complete mapping from the model parameters to the implied revenue ratio.

service prior to billing, as well as to provide simple cancellation procedures. Specifically, we simulate the effect of making consumers fully attentive every m months, with $m \in \{1, 3, 6, 12, 18, 24\}$.

Figure 5 reports the revenue ratio associated with each service, when "forced" attention varies from monthly (fully attentive) to these lower frequencies. It illustrates that the revenue impact is meaningful (relative to fully attentive subscribers) even with a 3-month frequency, but that, say, a 6-month frequency of "forced" attention would reduce the overall revenue impact of inattention by about 50%.

5.3 Robustness and heterogeneity

Robustness. In order to assess the robustness of the quantitative results to some of our modeling assumptions regarding inattention, we estimated several alternative specifications of the model. Panel A of Table 3 summarizes these results by reporting summary statistics for the estimates of λ and the revenue ratios, as well as how they correlate (across services) with the baseline model.

The first two rows report results from specifications in which we allow inattention to vary over time. In the first row of Panel A, we allow λ to vary linearly in time since subscription (that is, $\lambda_t = \lambda_0 + \theta t$), and in the second row we repeat the same model, but assume that λ "resets" to λ_0 after card replacement. The results from both specifications are remarkably similar to the baseline results.

The third and fourth rows of Panel A consider the possibility that when a card is replaced, some merchants may be able to update payment information automatically for some of their subscribers. If this is the case, some subscribers may not be fully attentive even at card expiration. To assess how sensitive the results are to this assumption, we re-estimate the baseline model, but assume that $\lambda = 0.75$ (third row) and $\lambda = 0.5$ (fourth row) at card expiration, instead of our baseline assumption (of $\lambda = 1$). Our estimates regarding λ remain almost the same, although the revenue ratio results are even more striking (because the revised assumptions affected the estimates of other model parameters).

Heterogeneity across consumers. Panel B of Table 3 summarizes the estimates of λ and the revenue ratio when we estimate the model separately for different subsets of consumers. We do not find meaningful heterogeneity when we estimate the model separately based on the demographics of the zip code where the account holder has the largest number of transactions (not reported).[24] This could reflect a lack of underlying correlation with demographics,

[24] When we compare estimates for the top quartile vs. the bottom quartile of zip-code level income, the resultant estimates of the two groups are essentially the same. There are similarly no differences based on

measurement error in zip code imputation, or selection patterns that make the demographics of subscribers less different than suggested by zip code demographics.

In contrast, we find meaningful and interesting heterogeneity at the card level, as reported in the first two rows of Panel B. In this exercise, we estimate the model separately for the 1.4% of subscribers in our sample who have used their card at some point for a cash advance, and compare the estimates to those for all other subscribers in our baseline sample (who have not used the card for a cash advance). Cash advances allow cardholders to borrow a certain amount of cash against their credit card's balance. It is often associated with fees and high interest rates, so is widely considered an expensive form of borrowing. Cash advances, therefore, may proxy for subscribers who are less financially sophisticated (Agarwal et al. 2009).

Appendix Figure A3 shows the retention rates for each service for cards with and without a cash advance, and it is evident that the drop in retention rates for cash-advance subscribers is sharper in the majority of the cases, which is indicative of greater inattention. In Panel B of Table 3, we show results from re-estimating the baseline model for these two groups of consumers. Consistent with the observed pattern of retention rates, the table shows that subscribers who have used cash advances in the past are less attentive (lower λ) and associated with a revenue ratio that is nearly twice as large on average. Moreover, this pattern seems to (weakly) persist across all services. Appendix Figure A4 plots the estimates of λ (left panel) and the revenue ratio (right panel) for cash-advance subscribers against other subscribers, service by service. Across all services cash-advance subscribers are associated with either the same or lower estimates of λ and a higher implied revenue ratio.

The use of cash advances is uncommon, so the above pattern is driven by a small fraction of subscribers. To address this, we construct a proxy for financial sophistication, which generates variation through the entire population of subscribers. Specifically, as detailed in Appendix B, we use cash advances as an outcome that indicates the lack of financial sophistication, and predict it using all other card transactions. This allows us to capture other aspects of card transactions, such as the amount of spending and the identity of the merchants the card holder is visiting, which may be indicative of financial sophistication. We use this predictive model to compute a predicted probability of using a cash advance for each subscriber in our sample, and we divide accounts into quartiles based on this predicted probability.

We re-estimate the model for each quartile and show the results in the bottom rows of Panel B of Table 3. We find that subscribers that are less financially sophisticated (that is, subscribers who are more likely to use a cash advance, denoted by Q4) tend to also be

zip code age and racial composition.

less attentive. Appendix Figure A5 illustrates the underlying patterns in the data: Subscribers who are less financially sophisticated tend to have a larger drop at the month of card replacement.

6 Conclusions

In this paper, we use payment card replacement as a way to explore the extent of inertia or inattention of consumers of subscription services. Using data from ten subscription services, we document a sharp drop in the rate of monthly renewal when the payment card is replaced, relative to other months when the subscription is automatically renewed without an active decision by the consumer.

Using a stylized model, we estimate that inattention and more passive renewal leads to a substantial increase in revenues for subscription services, which could be as high as three times greater than the revenues they would receive if subscribers were fully attentive and made an active renewal decision every month.

While this general result may be viewed as evidence of firms exploiting "behavioral" consumers in the spirit of DellaVigna and Malmendier (2004), it is important to note the convenience benefits associated with automatic renewal, which may make "forced" active renewal on a monthly basis a potentially undesirable remedy. We quantify the revenue impact of intermediate remedies, which require active renewal at lower frequencies. In some subscription situations, especially those associated with digital goods, one could also consider renewal notices that rely on account activity. Exploring these cases would require different data from that used in this study, and thus we leave it as an important topic for future research.

A key limitation of our study is that we condition on the initial subscription. As Miller, Sahni, and Strulov-Shlain (2023) document in the context of newspaper subscriptions, the propensity to subscribe is also affected by the auto-renewal features of the contract, and more broadly combining the two margins – initial subscription and monthly renewal – would be another promising direction for future work.

References

Agarwal, Sumit, John C Driscoll, Xavier Gabaix, and David Laibson. 2009. "The Age of Reason: Financial Decisions over the Life Cycle and Implications for Regulation." *Brookings Papers on Economic Activity* 2009 (2): 51–117.

Ascarza, Eva, Raghuram Iyengar, and Martin Schleicher. 2016. "The perils of proactive churn prevention using plan recommendations: Evidence from a field experiment." *Journal of Marketing Research* 53 (1): 46–60.

Chen, Chao, Andy Liaw, Leo Breiman, et al. 2004. "Using random forest to learn imbalanced data." *University of California, Berkeley* 110 (1-12): 24.

Consumer Financial Protection Bureau. 2021. *The Consumer Credit Card Market.* Technical report.

DellaVigna, Stefano, and Ulrike Malmendier. 2004. "Contract design and self-control: Theory and evidence." *The Quarterly Journal of Economics* 119 (2): 353–402.

———. 2006. "Paying Not to Go to the Gym." *The American Economic Review,* 694–719.

Di Geronimo, Linda, Larissa Braz, Enrico Fregnan, Fabio Palomba, and Alberto Bacchelli. 2020. "UI dark patterns and where to find them: a study on mobile applications and user perception." In *Proceedings of the 2020 CHI conference on human factors in computing systems,* 1–14.

Einav, Liran, Chiara Farronato, Jonathan Levin, and Neel Sundaresan. 2018. "Auctions versus posted prices in online markets." *Journal of Political Economy* 126 (1): 178–215.

Emarsys. 2021. *Everything Has Changed: The Rise of the Subscription and On-Demand Economies.* Available at https://emarsys.com/learn/blog/the-rise-of-the-subscription-and-on-demand-economies/. Last accessed Feb 9, 2023.

Experian. 2023. *What Is a Credit Card Cash Advance Fee?* Available at https://www.experian.com/blogs/ask-experian/what-is-credit-card-cash-advance-fee/. Last accessed June 15, 2023.

Federal Trade Commission. 2021. *Enforcement Policy Statement Regarding Negative Option Marketing.* Technical report.

———. 2023a. *FTC Takes Action Against Amazon for Enrolling Consumers in Amazon Prime Without Consent and Sabotaging Their Attempts to Cancel.* Available at https://www.ftc.gov/news-events/news/press-releases/2023/06/ftc-takes-action-against-amazon-enrolling-consumers-amazon-prime-without-consent-sabotaging-their. Last accessed July 28, 2023.

———. 2023b. *Negative Option Rule: Notice of Proposed Rulemaking - 88 FR 24716.* Available at https://www.federalregister.gov/documents/2023/04/24/2023-07035/negative-option-rule.

Johnen, Johannes. 2019. "Automatic-renewal contracts with heterogeneous consumer inertia." *Journal of Economics & Management Strategy* 28 (4): 765–786.

Luguri, Jamie, and Lior Jacob Strahilevitz. 2021. "Shining a light on dark patterns." *Journal of Legal Analysis* 13 (1): 43–109.

Mathur, Arunesh, Gunes Acar, Michael J Friedman, Eli Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. "Dark patterns at scale: Findings from a crawl of 11K shopping websites." *Proceedings of the ACM on Human-Computer Interaction* 3 (CSCW): 1–32.

Miller, Klaus M, Navdeep S Sahni, and Avner Strulov-Shlain. 2023. "Sophisticated Consumers with Inertia: Long-Term Implications from a Large-Scale Field Experiment." Available at SSRN: https://ssrn.com/abstract=4065098 or https://dx.doi.org/10.2139/ssrn.4065098. Last accessed Jun 7, 2023.

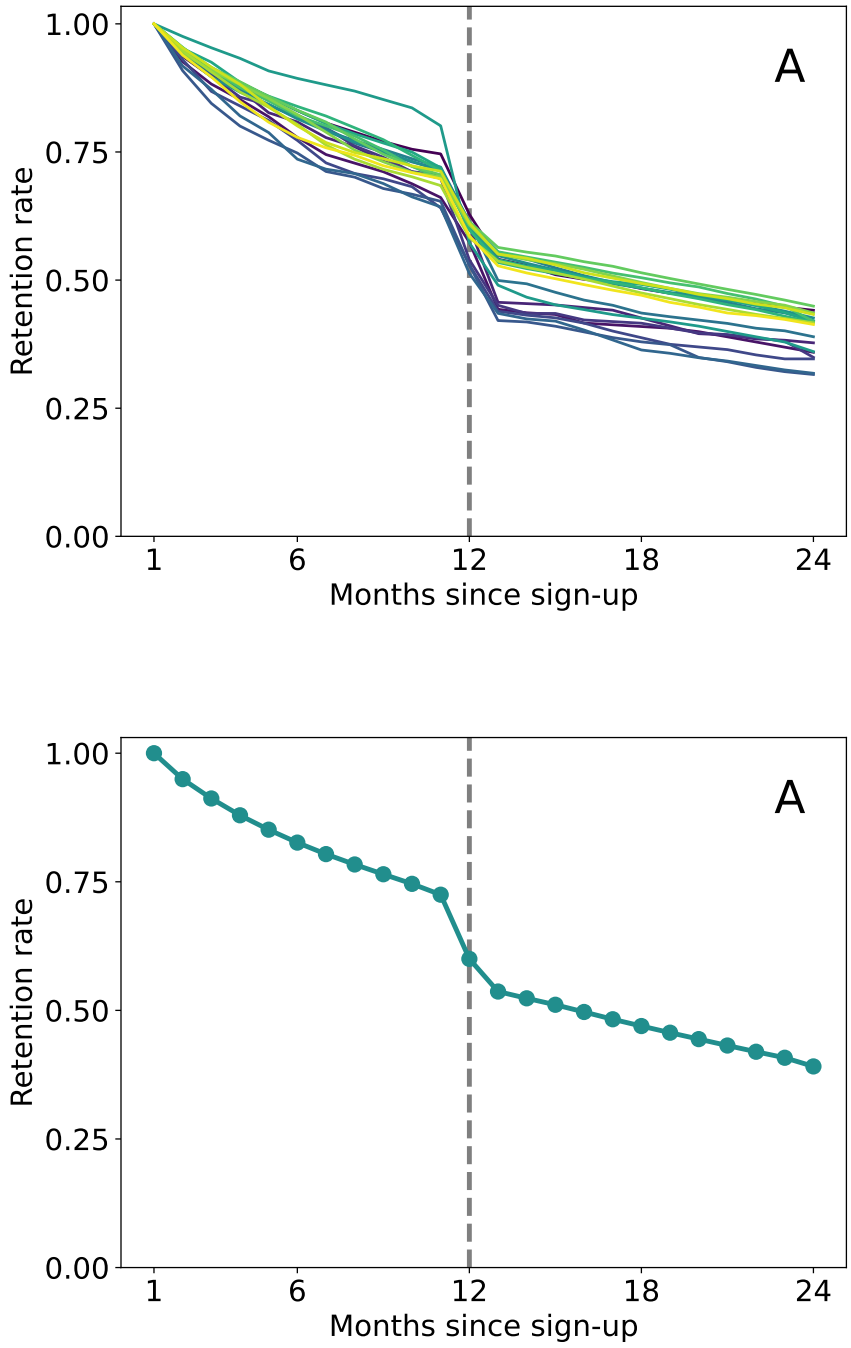
Posner, Nathaniel, Andrey Simonov, Kellen Mrkva, and Eric J Johnson. 2022. "Dark Defaults: How Choice Architecture Steers Campaign Donations." *Available at SSRN 4258478.*

Reme, Bjørn-Atle, Helene Lie Røhr, and Morten Sæthre. 2022. "Inattention in Contract Markets: Evidence from a Consolidation of Options in Telecom." *Management Science* 68 (2): 1019–1038.

West Monroe. 2021. *The State of Subscription Services Spending.* Technical report.

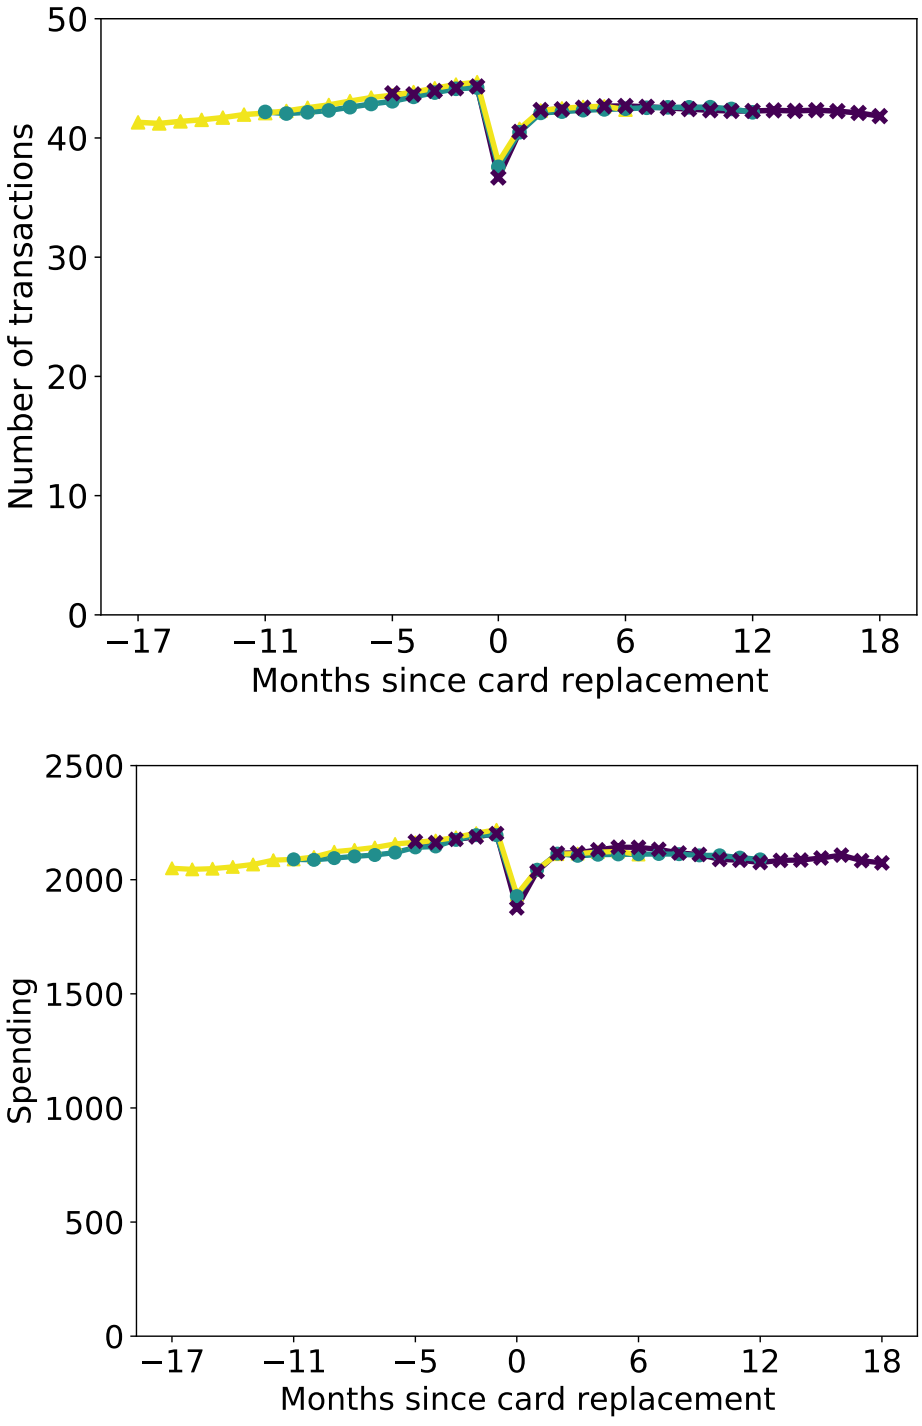
Zuora. 2022. *The Subscription Economy Index.* Technical report. Available at https://www.zuora.com/resources/subscription-economy-index/. Last accessed Feb 9, 2023.

Figure 1: Aggregating retention rates



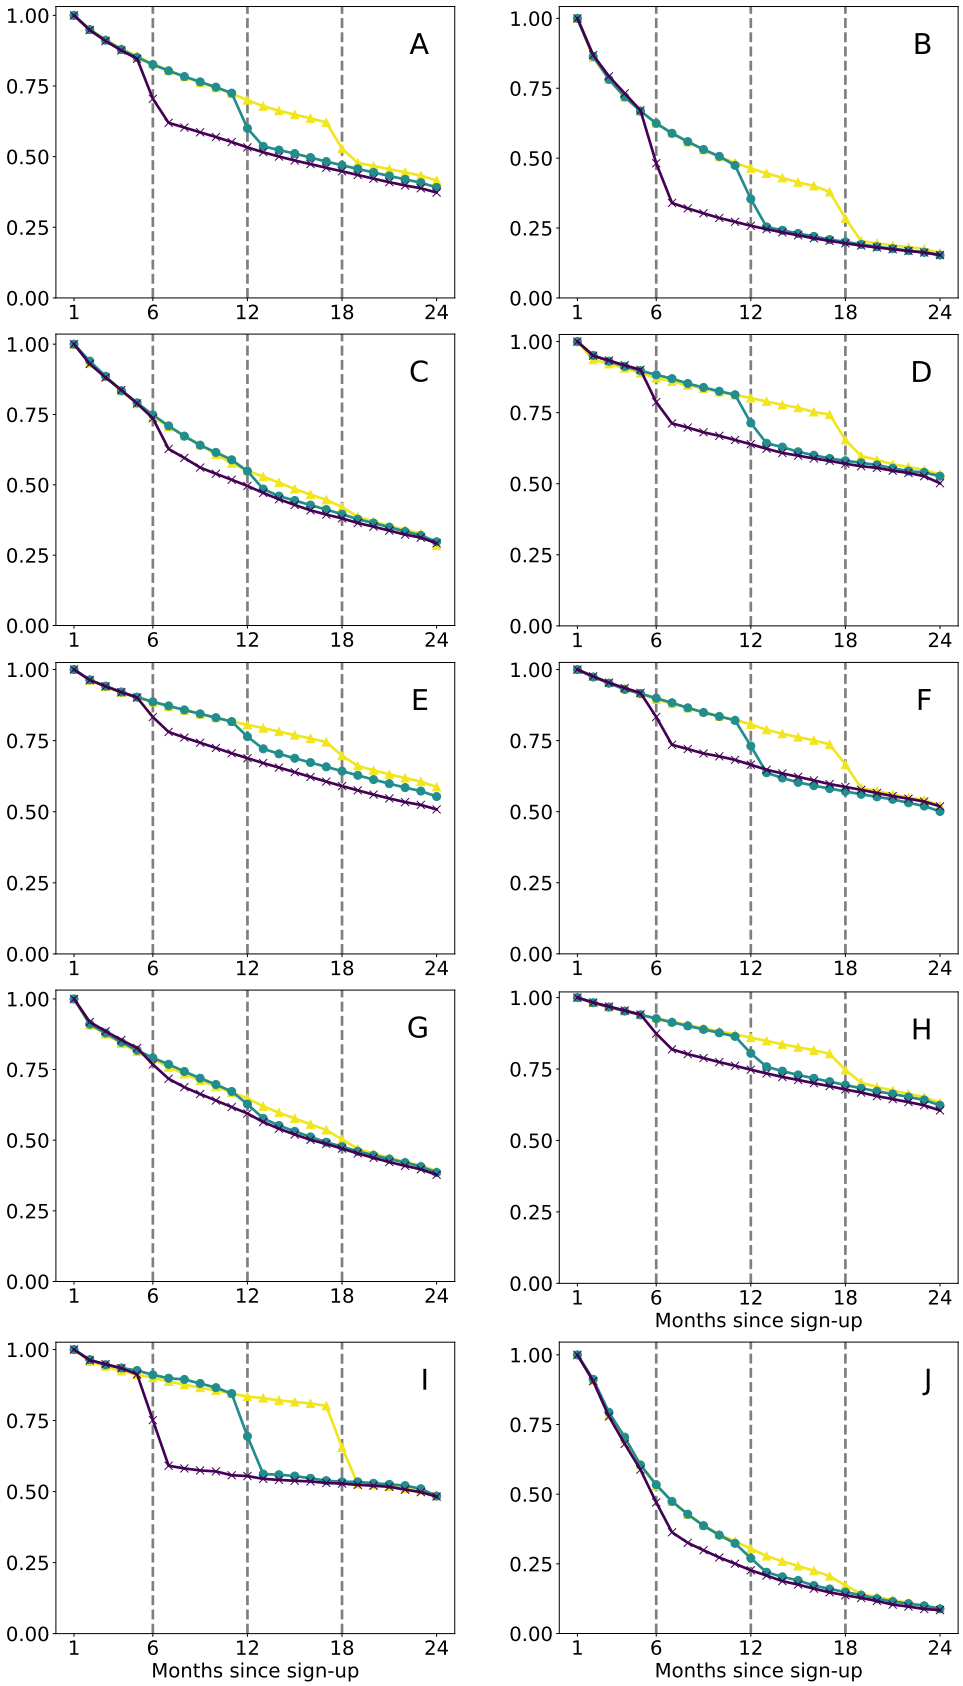
Note: Figure shows retention rates by month, for subscription service A and cohorts with card replacement 12 months after the initial subscription. The top panel shows the raw retention rates for all 19 cohorts. The bottom panel shows the adjusted retention rate $\widehat{R}_n(x)$, which aggregates across cohorts using equations (2) and (3) (see Section 3 for more details).

Figure 2: Account activity around replacement date



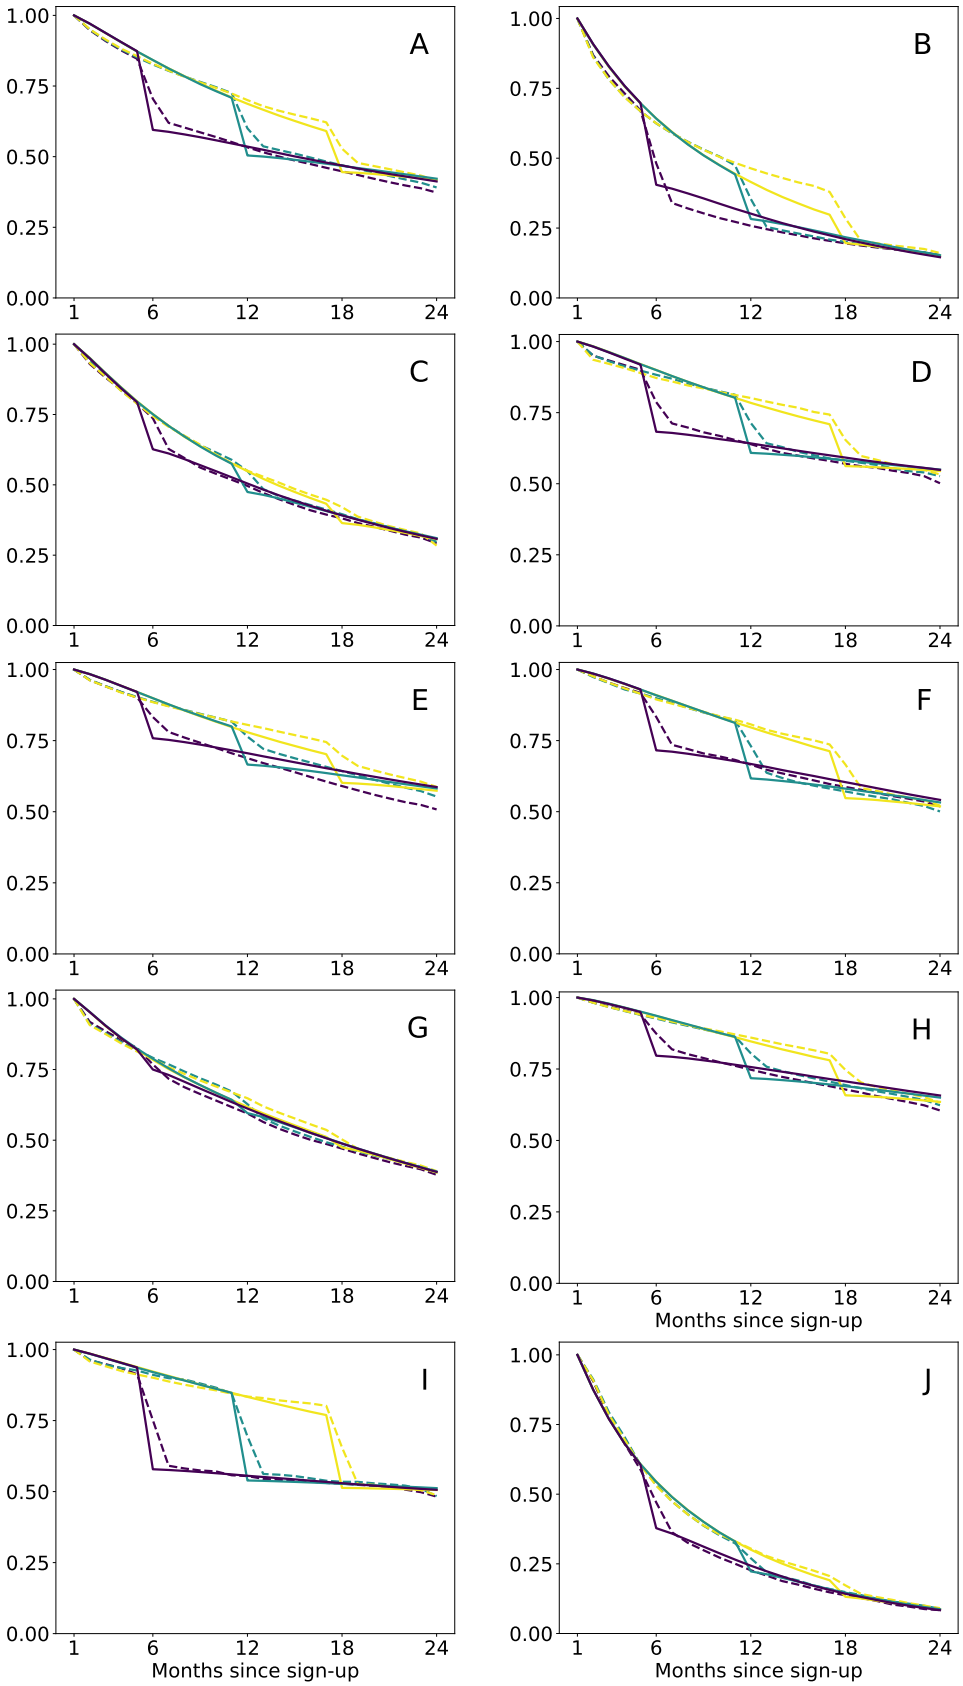
Note: Figure shows account activity, measured by number of transactions (top panel) and spending (bottom panel) per month, around the month of card replacement. To construct the figure, we calculate transactions and spending in each month, including transactions at the ten subscription services that we study. We regress this on account and month fixed effects and average the residuals by months from the date of card replacement. The plot shows the average of the residuals plus the mean number of transactions and spending per month computed across the entire sample for our three groups of cohorts (accounts that subscribe to one of the ten services 6, 12, or 18 months prior to the date of card expiration.)

Figure 3: Retention rate by month since initial subscription, all services and cohorts



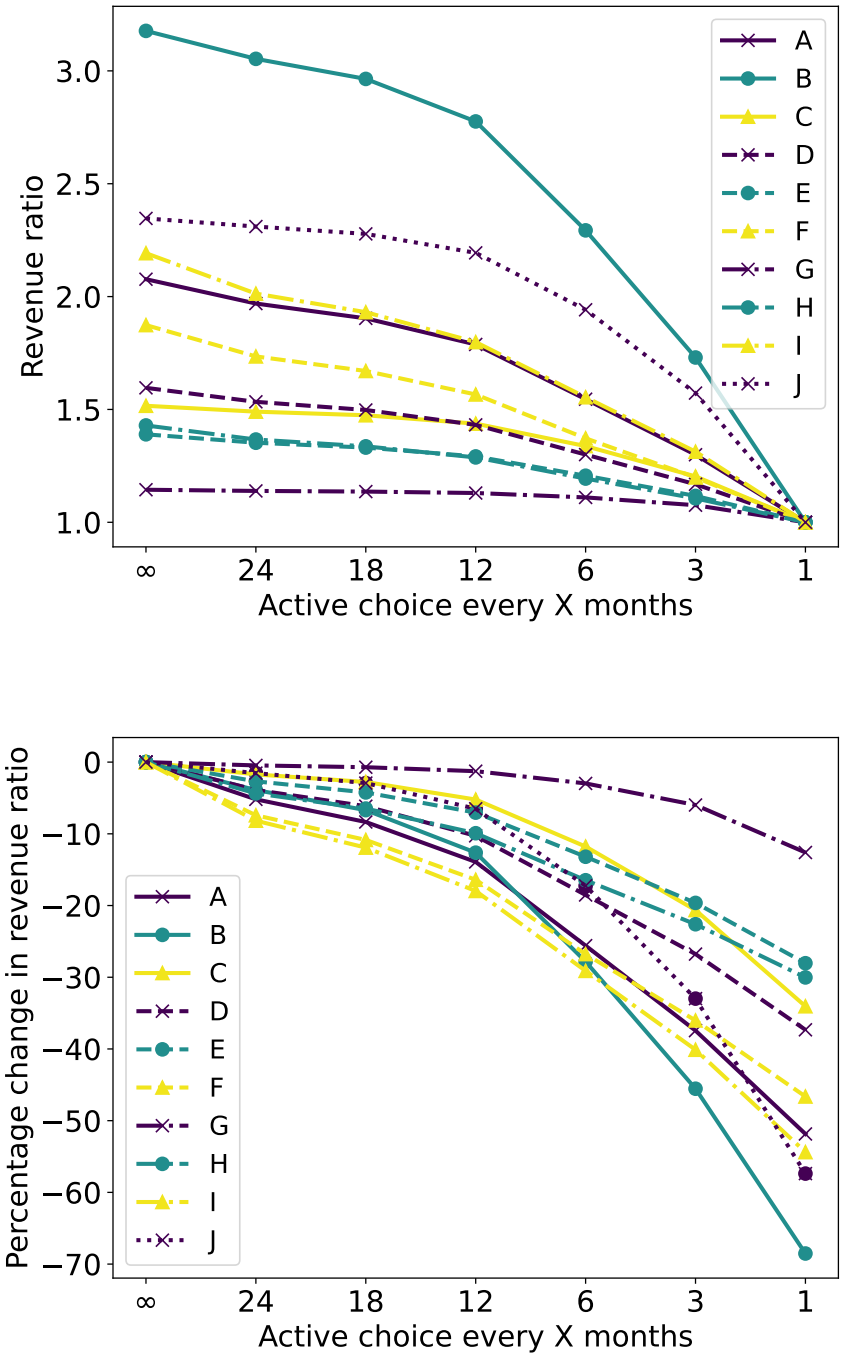
Note: Figure shows, for all subscription services, the aggregated and seasonally adjusted retention rate $\widehat{R}_n(x)$ by month since sign-up for the three groups of cohorts with card replacement 6, 12, or 18 months after initial sign-up. The letter in the top right corner of each panel identifies the subscription service. For each service, we adjust for seasonality by removing calendar month-year fixed effects common to all cohorts. We describe the procedure for seasonal adjustment in detail in Section 2.

Figure 4: Model fit, all services and cohorts



Note: Figure shows, for all subscription services, the retention rate $\widehat{R}_n(x)$ by month since sign-up (dashed lines) as in Figure 3, and the retention rates (solid lines) as predicted by our model. The letter in the top right corner of each panel identifies the subscription service.

Figure 5: Revenue impact of possible regulatory remedies



Note: Figure shows, for all subscription services, the revenue impact of forcing subscribers to make an active choice every 1, 3, 6, 12, 18, 24 months or never (∞). The top panel shows the revenue under active choice every X months relative to the revenue if subscribers were attentive every month/are forced to make an active choice every month. The bottom panel shows the percentage change in this ratio relative to the default of never being forced to make an active choice (∞). We construct the revenue ratio as follows: For each service, we simulate the monthly subscription choice of 100,000 hypothetical subscribers and for 120 months after initial sign-up. The denominator is the discounted sum of monthly subscribers if they were forced to make an active choice every month. The numerator is the discounted sum of monthly subscribers if they were forced to make an active choice every X months. The default case is that subscribers are never forced to make an active choice (∞) and only pay attention with probability λ every month. Appendix Table A2 provides the underlying numbers associated with the top panel.

Table 1: Parameter estimates by service

Service	ρ	λ	η
A	0.990	0.110	0.823
	(0.003)	(0.002)	(0.021)
B	0.856	0.184	0.004
	(0.009)	(0.002)	(0.000)
C	0.946	0.270	1.588
	(0.006)	(0.014)	(0.106)
D	1.019	0.090	1.404
	(0.014)	(0.006)	(0.133)
E	0.999	0.133	2.558
	(0.002)	(0.002)	(0.047)
F	0.978	0.093	1.969
	(0.008)	(0.004)	(0.123)
G	0.946	0.502	3.784
	(0.002)	(0.000)	(0.074)
H	0.995	0.097	3.282
	(0.003)	(0.002)	(0.080)
I	1.053	0.044	0.477
	(0.040)	(0.003)	(0.105)
J	0.819	0.277	0.154
	(0.011)	(0.021)	(0.113)

Note: Table reports parameter estimates and bootstrap standard errors (in parentheses) for our model of subscription renewal behavior described in Section 4. We estimate the model separately for each of the 10 subscription services A through J. To compute the standard errors, we estimate our model on 1,000 bootstrap samples. For each parameter estimate, the standard error is the standard deviation of the 1,000 bootstrap estimates.

Table 2: Revenue impact of inattention

Service	Share unaffected	Avg months subscribed		Revenue ratio
		If inattentive	if attentive	
A	0.05	36.6	13.2	2.08
B	0.00	14.1	4.1	3.18
C	0.00	21.6	13.6	1.52
D	0.27	41.7	9.5	1.60
E	0.21	41.2	20.2	1.39
F	0.04	44.7	19.9	1.87
G	0.00	23.8	20.4	1.14
H	0.20	49.6	24.4	1.43
I	0.25	54.1	4.1	2.19
J	0.00	10.0	4.1	2.35
Mean	**0.10**	**33.7**	**13.4**	**1.87**

Note: Table reports our counterfactual estimates on how inattention affects revenue, separately for the 10 subscription services A through J. For each service, we simulate the monthly subscription choice of 100,000 hypothetical subscribers for 120 months after initial sign-up. Column (1) reports the share of subscribers not affected by inattention because they have a positive valuation in every month. Columns (2) and (3) show, for subscribers with a negative valuation in at least one month, the average number of months they are subscribed if inattentive (paying attention with probability λ each month) and if attentive (paying attention with probability 1 each month), respectively. Column (4) shows the gain in relative revenue from inattention over both affected and unaffected subscribers. We construct the revenue ratio as follows. First, we compute the number of subscribers in each month for the case that subscribers pay attention with probability λ. We sum these numbers across month, discounting at a rate of 1%. We divide the discounted sum of monthly subscribers if inattentive by the discounted sum of monthly subscribers if subscribers were fully attentive (paying attention with probability 1).

Table 3: Robustness and heterogeneity

	Lambda					Revenue ratio				
	Mean	St. Dev.	2nd	9th	Corr. w/ baseline	Mean	St. Dev.	2nd	9th	Corr. w/ baseline
Baseline	0.18	0.13	0.09	0.28		1.87	0.57	1.39	2.35	
A. Robustness										
Linear decay of lambda	0.16	0.09	0.09	0.28	0.93	2.88	3.43	1.27	2.67	0.84
Linear decay of lambda with reset	0.17	0.12	0.09	0.28	1.00	2.65	1.88	1.37	2.89	0.86
Lambda at card expiration = 0.75	0.17	0.13	0.07	0.26	0.99	2.55	1.44	1.46	3.57	0.63
Lambda at card expiration = 0.5	0.18	0.15	0.04	0.37	0.97	4.13	4.34	1.53	5.46	0.32
B. Heterogeneity										
Never used cash advance	0.17	0.11	0.09	0.30		1.84	0.56	1.35	2.18	
Used cash advance	0.14	0.11	0.05	0.22		3.42	2.55	1.43	5.65	
Financial sophistication -- Q1	0.20	0.11	0.05	0.36		2.04	1.21	1.27	3.51	
Financial sophistication -- Q2	0.18	0.12	0.08	0.34		1.84	0.67	1.26	2.19	
Financial sophistication -- Q3	0.17	0.12	0.08	0.31		1.89	0.56	1.49	2.28	
Financial sophistication -- Q4	0.18	0.12	0.08	0.30		1.93	0.54	1.41	2.26	

Note: Table summarizes the results of our robustness and heterogeneity analyses for the estimates of λ and the revenue ratio. The first row summarizes our baseline estimates. We report the mean and standard deviation of the respective estimates across the ten services, as well as the second and ninth value of the estimates if sorted in ascending order (which corresponds to the 10th and 90th percentiles). Panel A summarizes the results of our robustness analysis, where we modeled alternative specifications regarding inattention. The first row summarizes the results of a specification that allows λ to vary linearly in time since subscription, that is $\lambda_t = \lambda_0 + \theta t$. In the second row, the model again allows λ to vary linearly, but we assume that λ_t "resets" to λ_0 after card expiration. That is, $\lambda_t = \lambda_0 + \theta t$ for $t < x$, and $\lambda_t = \lambda_0 + \theta(t - x)$ for $t > x$, where x is the month of card expiration. For λ, we compute the average "experienced" λ for each service. That is, we weight the period-specific λ_t by the share of consumers still subscribed in period t, normalized so that weights add up to one. We do so using the observed, regression-adjusted retention rates of each cohort, omitting the period of card expiration. We average across the three cohorts of card expiration, weighing by the total number of initial subscribers. In the third and fourth row, we estimate our baseline model with time-invariant λ, but assume that $\lambda = 0.75$ and $\lambda = 0.5$ in the month of card expiration, respectively (instead of $\lambda = 1$). We report the mean and standard deviation of the estimates of the ("experienced") λ and revenue ratio, the second and ninth value, as well as the correlation with the baseline estimates. Panel B summarizes the results of our heterogeneity analysis, where we estimate our baseline model separately for subsets of cards. First, we split the sample of cards by whether they ever had a cash advance. Second, we split the sample of cards into quartiles based on their financial sophistication, which we measure by the predicted probability of using a cash advance.
We measure financial sophistication as the predicted probability that a card has a cash advance. Appendix B describes how we predict this probability and Appendix Table A3 shows descriptive statistics on accounts grouped by quartiles of the predicted probability. We report the estimation results for all quartiles of financial sophistication: Q1 corresponds to the lowest predicted probability of having a cash advance (that is, highest level of financial sophistication) and Q4 corresponds to the highest predicted probability of having a cash advance (that is, lowest level of financial sophistication).

Online Appendix

Appendix A: Simulating model predictions

Let $\varphi = (\rho, \lambda, \eta)$ be the set of candidate parameter values. In order to construct model predictions for retention rates as a function of these parameters, we simulate a large panel of initial subscribers and record their renewal decisions as given by the model.

For each subscription service, and for each number of months between initial subscription and card replacement $x = 6, 12, 18$, we simulate three sets of $N = 100,000$ subscribers, and then simulate their renewal decisions as a function of the model and the parameters given by φ.

Specifically, we start by drawing the random components of the model, which include: (i) an $N \times 1$ vector of draws that will affect initial valuations – iid draws from a $[0, 1]$-uniform distribution – denoted by $u_0(x)$; (ii) an $N \times T$ matrix of taste shocks – iid draws from a standard normal distribution – denoted by $e_0(x)$; and (iii) an $N \times T$ matrix of "attention shocks" – iid draws from a $[0, 1]$-uniform distribution – denoted by $l_0(x)$.

We then construct the vector of initial values, $v_0(x)$, by transforming u_0 into a variable with exponential distribution and mean η using $v_0(x) = -\eta \log(1 - u_0(x))$. $v_0(x)$ is then an $N \times 1$ vector of initial net flow utilities in month $s = 1$. Then, using $v_0(x)$ and $e_0(x)$, we simulate the net flow utility for all subsequent periods using ρ and the model assumption that $v_t(x) = \rho v_{t-1}(x) + e_0(x)$.

We then say that subscriber i pays attention in month s if and only if $l_0(x) \leq \lambda$. Denote by $a(x)$ the resulting $N \times T$ matrix of binary attention indicators, and we assign $a(x) = 1$ for month $s = x$ for all subscribers.

Finally, we construct the $N \times T$ subscription matrix, $s(x)$. As all individuals are subscribers in $s = 1$, we have that $s(x)_{i1} = 1$ for all i. For $t > 1$, $s(x)_{it} = 0$ if $s(x)_{i,t-1} = 0$ or if $a(x)_{it} = 1$ and $v(x)_{it} < 0$. The simulated retention rates for each month s are then given by the share of 1s in each column of the subscription matrix $s(x)$, which we denote by $R(x; \varphi)$.

Appendix B: Predicting cash advance

We do additional analysis on the subset of cards that took a cash advance. A cash advance is a withdrawal of cash against a cardholder's credit limit and is typically associated with high fees and interest rates.[25] These transactions are identified separately in the payment card data as cash withdrawals from a credit account (we do not consider cash withdrawals from debit accounts as these typically are charged only an ATM fee). 1.4% of accounts in our analysis sample took at least one cash advance during our sample period.[26]

[25] Fees and interest rates vary by bank and account type. Typical charges are the higher of 5% of the value of the withdrawal or $10, plus an interest rate that is higher than that used for purchases, and which begins to accrue immediately with no grace period (Experian 2023).

[26] This is in line with survey evidence from Consumer Financial Protection Bureau (2021) which reports that approximately 2% of all consumers used a cash advance in 2021.

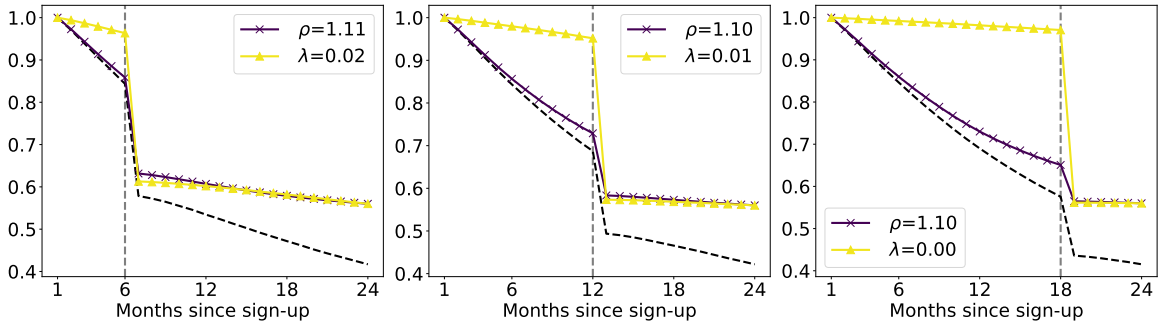
Because relatively few consumers take cash advances, the sample sizes for this subset are small. To remedy this, we fit a predictive model of the propensity to use cash advance from each cardholder's purchases. We use this model to predict the probability that each cardholder takes a cash advance.

Specifically, for each cardholder in our analysis sample, we create explanatory variables using the account's transaction history during the 6 months prior to card replacement (i.e. from months -6 to months -1). As predictors, we use the average number of monthly transactions, average monthly spending, share of online transactions, and the share of transactions at each of the top 50 nationwide merchants. We use this set of predictors to fit a random forest classifier to predict the probability of using cash advance. Because our data is imbalanced (i.e. many more cards do not use cash advance than do), we follow Chen, Liaw, Breiman, et al. (2004), who propose over-sampling the minority group. We use a model with 500 trees and a maximum depth of 50 nodes, and estimate it using five-fold cross validation. We then use the fitted classifier to predict the probability that a card used a cash advance for each card in the sample.

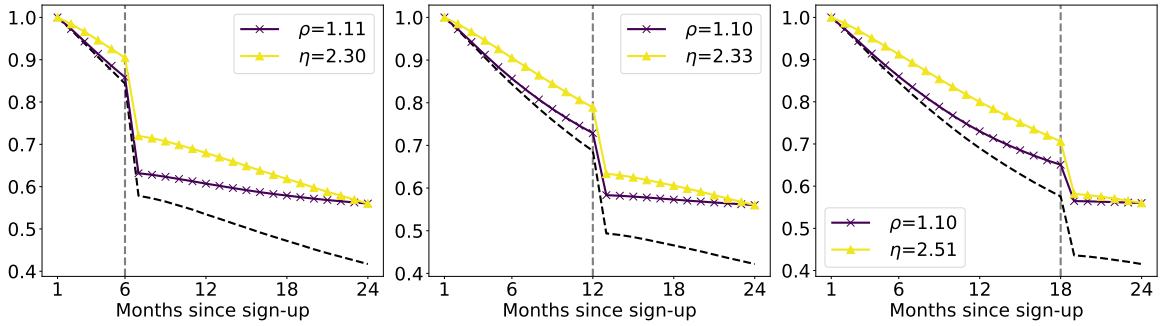
Using this predicted probability of cash advance, we split subscribers into quartiles. Appendix Table A3 shows descriptive statistics on accounts by quartile, where quartile 1 has the lowest predicted probability and quartile 4 has the highest. We then construct survival curves for accounts in each of these quartiles, which we show in Appendix Figure A5.

Appendix Figure A1: Comparative statics

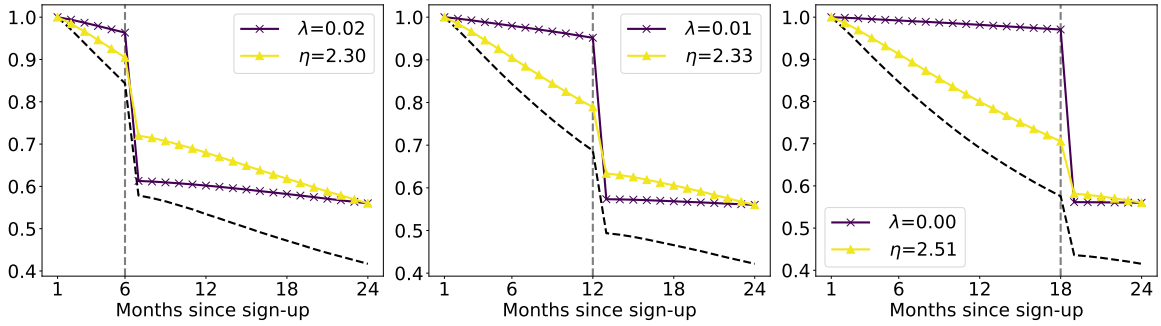
Holding $\eta = 0.855$ fixed:



Holding $\lambda = 0.111$ fixed:

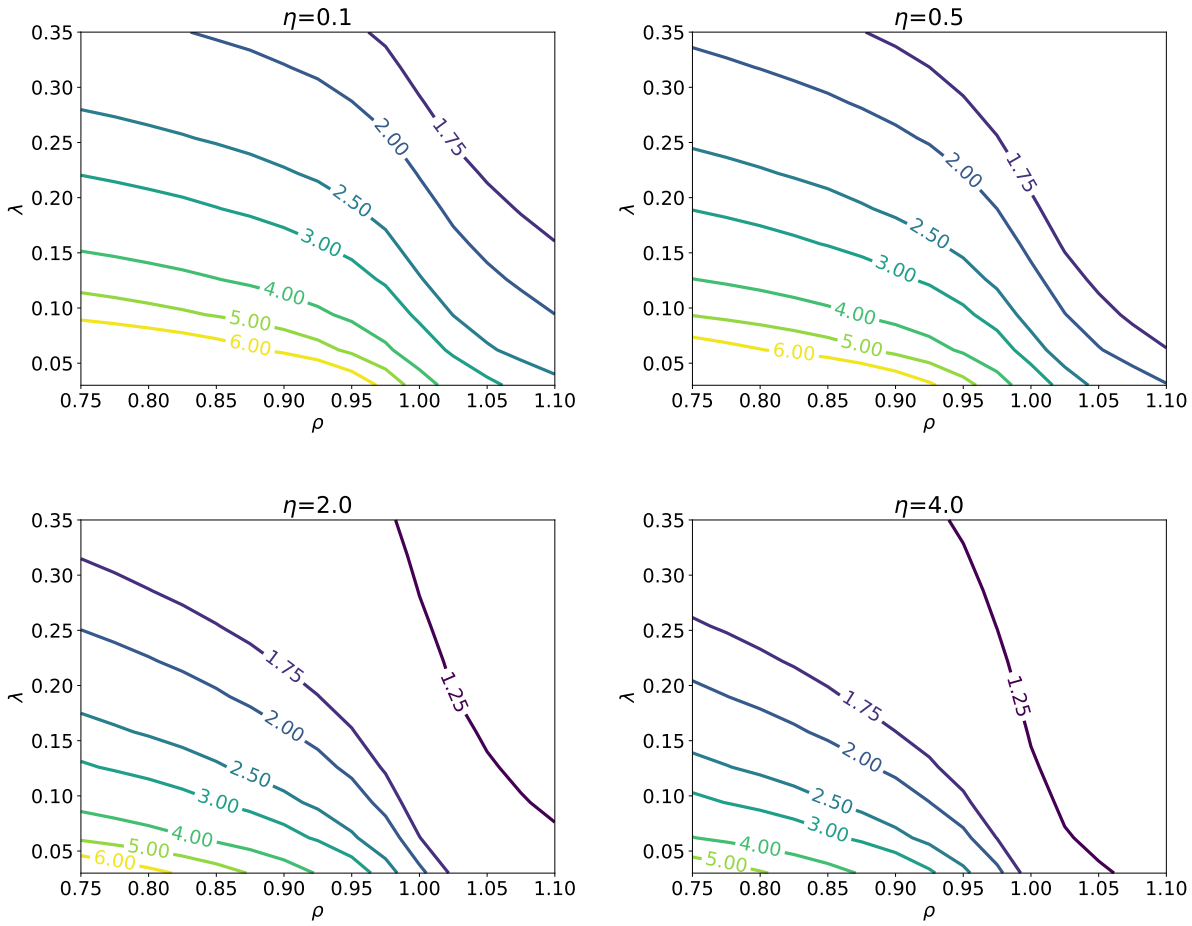


Holding $\rho = 0.988$ fixed:



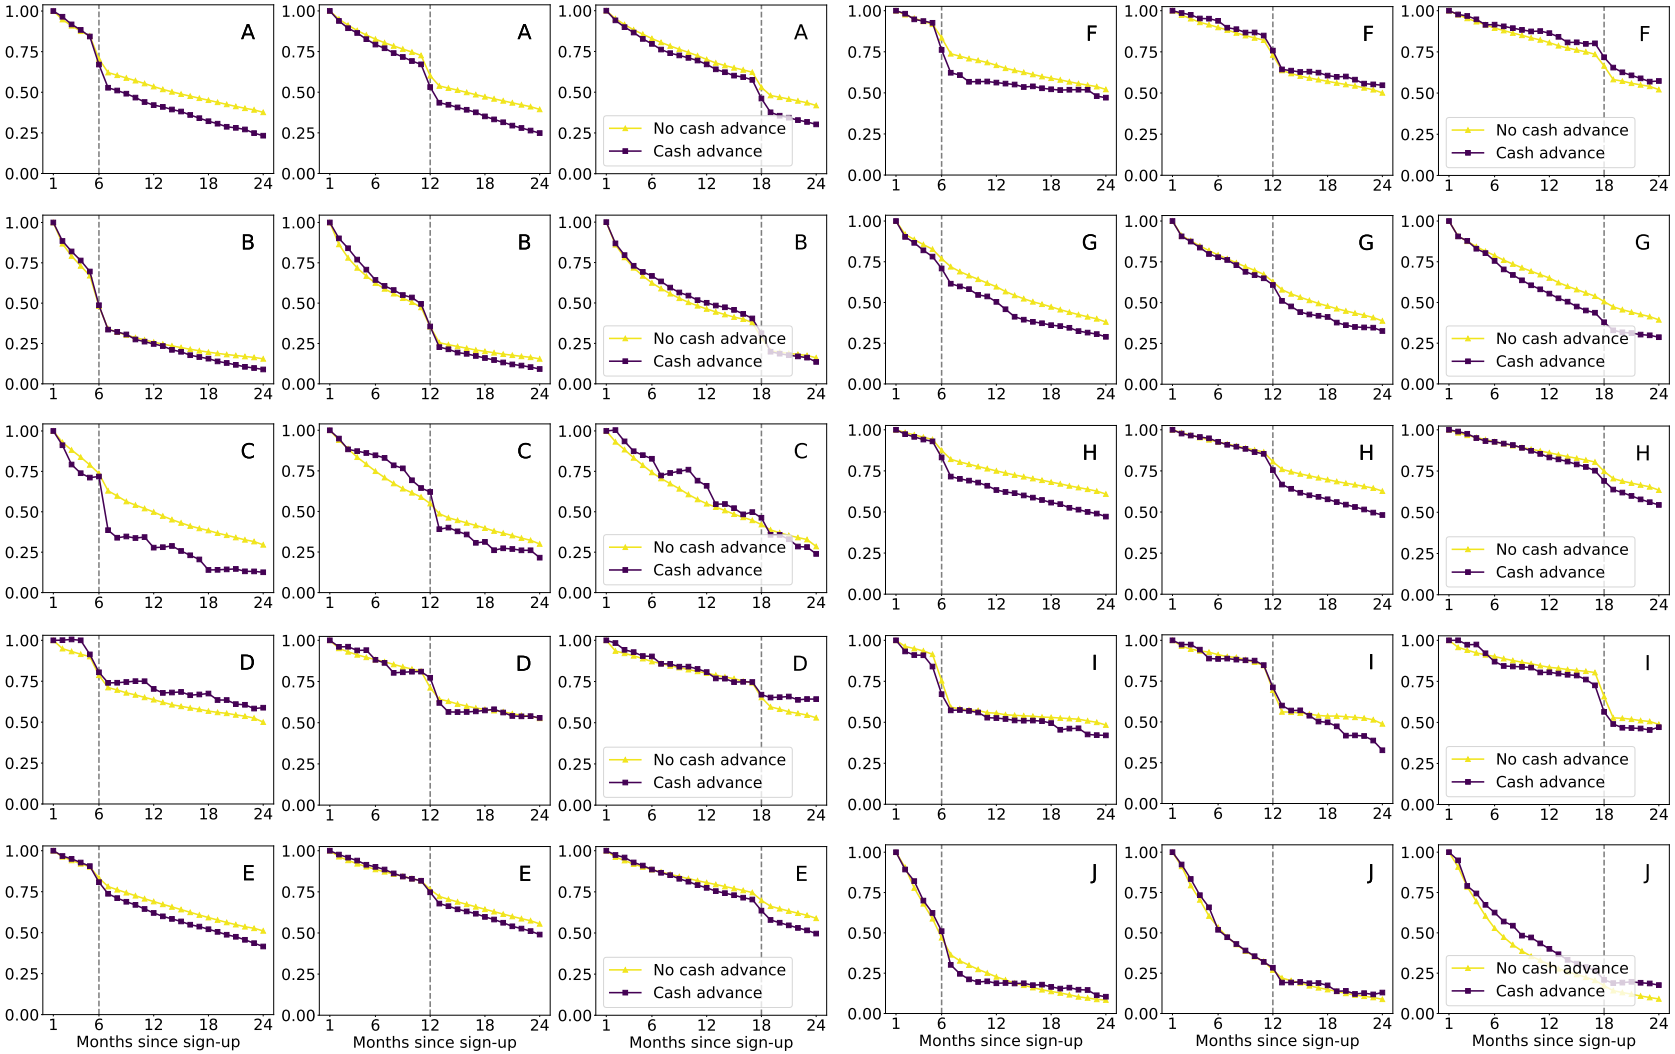
Note: Figure shows the comparative statics of our model of subscription renewal behavior using the data and estimated parameters for service A. Each panel shows model-predicted retention rates for service A on the y-axis, by months since sign-up on the x-axis. We consider the three groups of cohorts separately, with card replacements after 6 months (left panels), 12 months (middle panels), and 18 months (right panels). In each panel, the vertical line represents the month of card replacement and the black dashed line shows the model-predicted retention rates implied by our parameter estimates $\rho = 0.988$, $\lambda = 0.111$, and $\eta = 0.855$ (see Table 1). In each panel, the colored, solid lines illustrate the trade-off between pairs of model parameters, (ρ, λ), (ρ, η), and (λ, η), in matching the (arbitrary) target retention rate of 0.56 in the last month. In the top row, we keep $\eta = 0.855$ fixed and show, separately for each group of cohorts, how either ρ (also holding $\lambda = 0.111$ fixed) or λ (also holding $\rho = 0.988$ fixed) has to change to attain the target retention rate. For example, in the top-left panel, we attain this retention rate for the cohort of card replacements after 6 months by increasing ρ to $\rho = 1.11$, while holding $\lambda = 0.111$ and $\eta = 0.855$ fixed, or by decreasing λ to $\lambda = 0.02$, while holding $\rho = 0.988$ and $\eta = 0.855$ fixed. In the middle row, we hold $\lambda = 0.111$ fixed and show the trade-off between ρ and η. In the bottom row, we hold $\rho = 0.988$ fixed and show the trade-off between λ and η.

Appendix Figure A2: Iso-ratio curves



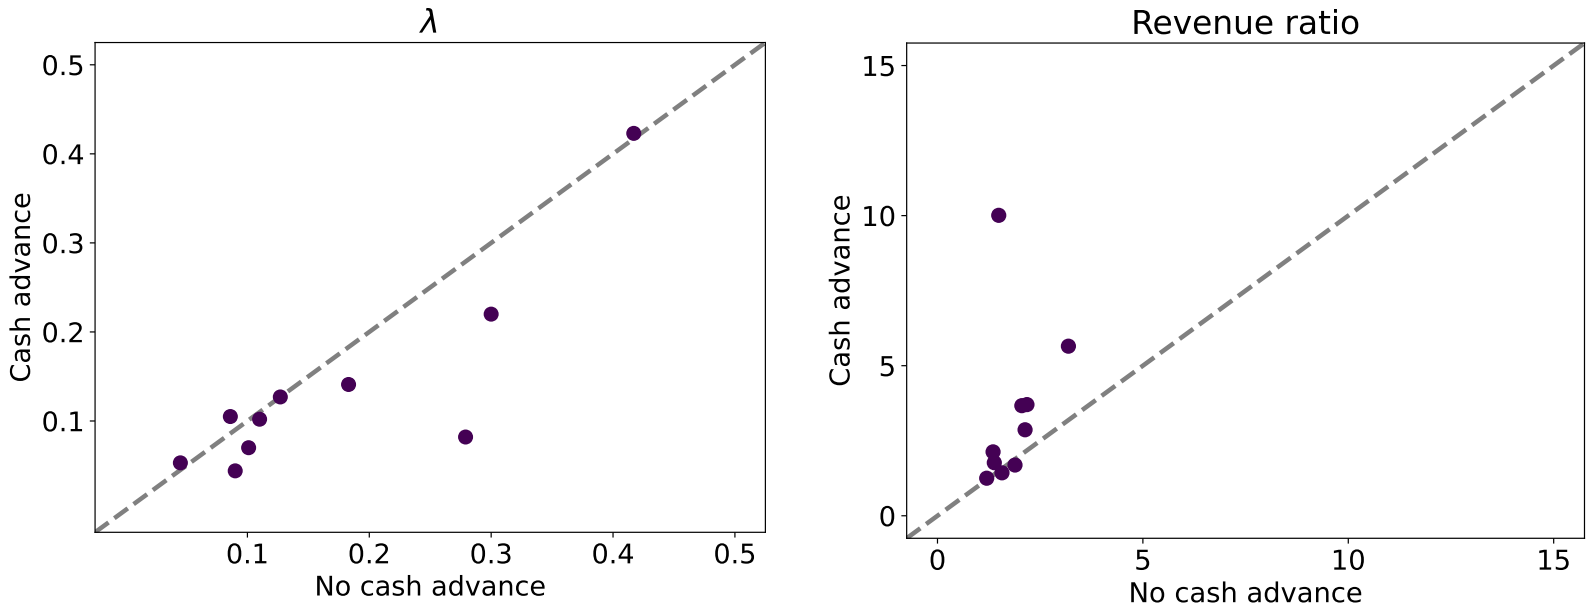
Note: Figure shows the revenue ratio as a function of the three model parameters (ρ, λ) and η. Given η, each line shows all (ρ, λ) combinations that yield the same revenue ratio. For given parameter values, we construct the revenue ratio by simulating the monthly subscription behavior of 100,000 hypothetical subscribers for 120 months after initial sign-up. The denominator of the ratio is the discounted sum of monthly subscribers if paying attention with probability 1 each month. The numerator is the discounted sum of monthly subscribers if inattentive, i.e., paying attention with probability λ in each month. We discount future revenues at a rate of 1%.

Appendix Figure A3: Heterogeneity by cash advance



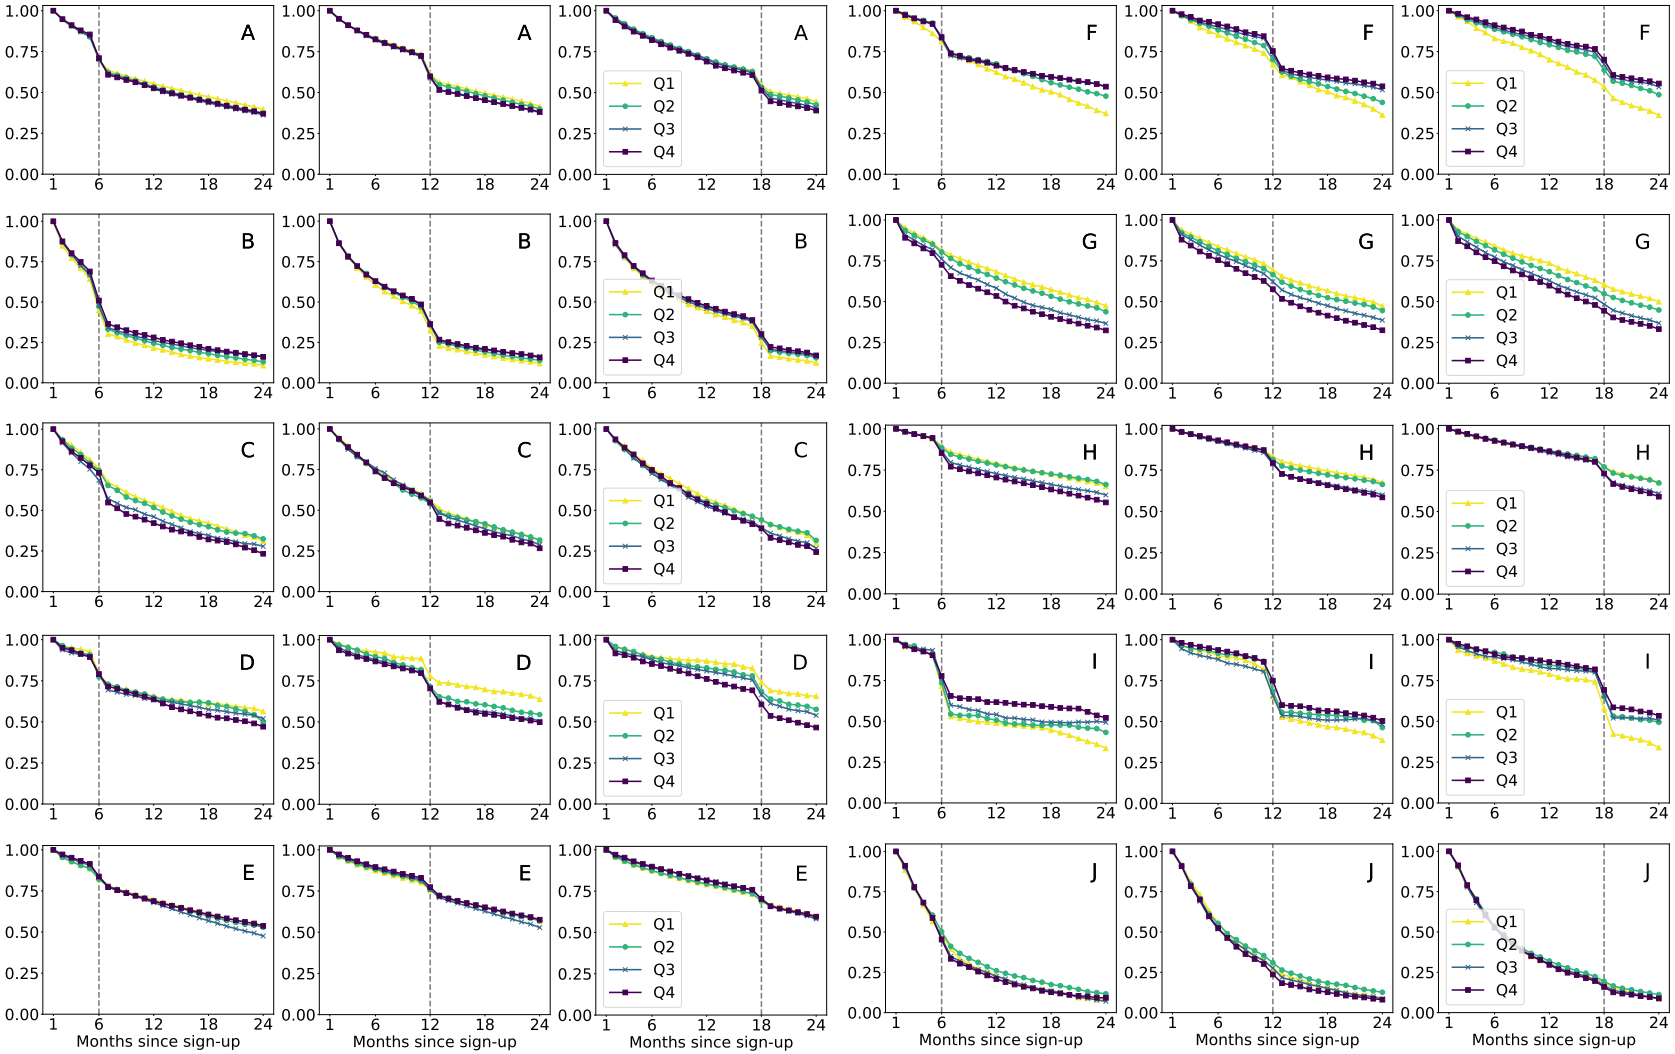
Note: Figure shows the regression-adjusted retention rates by month since sign-up for cards with and without cash advance and by group of cohorts with card replacement 6, 12, or 18 months after sign-up. The letter at the top right corner of each panel identifies the subscription service.

Appendix Figure A4: Estimation results by cash advance



Note: Figure illustrates the estimation results by cash advance for λ (left) and the revenue ratio (right) for each subscription service. In the figure on the left-hand side, we plot the estimate of λ for a given service and cards with cash advance (y-axis) against the estimate for cards without cash advance (x-axis). The figure on the right-hand side plots the revenue ratio for cards with cash advance against no cash advance. In each figure, the dashed line represents the 45-degree line.

Appendix Figure A5: Heterogeneity by financial sophistication



Note: Figure shows the regression-adjusted retention rates by month since sign-up, separately by quartile of financial sophistication and by group of cohorts with card replacement 6, 12, or 18 months after sign-up. We measure a cardholder's financial sophistication as one minus the predicted probability of taking a cash advance. Section 6 describes our procedure for predicting the probability that a cardholder takes a cash advance. We group accounts into quartiles based on the predicted probability, where Q1 captures the lowest probability of a cash advance, i.e., highest level of financial sophistication, and Q4 captures the highest predicted probability of a cash advance, i.e., lowest level of financial sophistication. We compute the regression-adjusted retention rates separately for each quartile. The letter at the top right corner of each panel identifies the subscription service.

Appendix Table A1: Average change in retention rates

	Average monthly change in retention rate	
Service	Replacement months	Other months
A	-0.09	-0.02
B	-0.12	-0.03
C	-0.05	-0.03
D	-0.08	-0.01
E	-0.05	-0.02
F	-0.09	-0.01
G	-0.05	-0.03
H	-0.05	-0.01
I	-0.15	-0.01
J	-0.07	-0.04
Mean	-0.08	-0.02

Note: Table shows, for each subscription service, the average monthly change in retention rates during months of card replacement and other months. For each service and cohort ($x = 6, 12, 18$), we compute the change in regression-adjusted retention rates, $\widehat{R}_n(x) - \widehat{R}_{n-1}(x)$. For each service, we average the change in retention rates across months and cohorts, separately for months with and without card replacement. The last line shows the average monthly change in retention rates averaged across the ten subscription services. For cohorts with card replacement six months after sign-up ($x = 6$), card replacement happens in months $n = 6$ and $n = 7$. For cohorts with card replacement 12 months after sign-up ($x = 12$), card replacement happens in months $n = 12$ and $n = 13$. For cohorts with card replacement 18 months after sign-up ($x = 18$), card replacement happens in months $n = 18$ and $n = 19$.

Appendix Table A2: Revenue impact of possible regulatory remedies

	Active choice every X months						
Service	∞	24	18	12	6	3	1
A	2.08	1.97	1.90	1.79	1.55	1.30	1.00
B	3.18	3.05	2.96	2.78	2.29	1.73	1.00
C	1.52	1.49	1.47	1.44	1.34	1.20	1.00
D	1.60	1.53	1.50	1.43	1.30	1.17	1.00
E	1.39	1.35	1.33	1.29	1.21	1.12	1.00
F	1.87	1.73	1.67	1.57	1.37	1.20	1.00
G	1.14	1.14	1.14	1.13	1.11	1.08	1.00
H	1.43	1.37	1.34	1.29	1.19	1.11	1.00
I	2.19	2.01	1.93	1.80	1.55	1.31	1.00
J	2.35	2.31	2.28	2.19	1.94	1.57	1.00
Mean	1.87	1.80	1.75	1.67	1.49	1.28	1.00

Note: Table summarizes the revenue impact of possible regulatory remedies as shown in the top panel of Figure 5. That is, the revenue impact of forcing subscribers to make an active choice every 1, 3, 6, 12, 18, 24 months or never (∞). Each column shows, for all subscription services, the revenue under active choice every X months relative to the revenue if subscribers were attentive every month/are forced to make an active choice every month. We construct the revenue ratio as follows: For each service, we simulate the monthly subscription choice of 100,000 hypothetical subscribers and for 120 months after initial sign-up. The denominator is the discounted sum of monthly subscribers if they were forced to make an active choice every month. The numerator is the discounted sum of monthly subscribers if they were forced to make an active choice every X months. The default case, as illustrated in Table 2, is that subscribers are never forced to make an active choice (∞) and only pay attention with probability λ every month.

Appendix Table A3: Quartiles of predicted cash advance

Predicted Pr(cash advance)		Share w/ cash advance	Mean monthly number of transactions	Mean monthly spend (USD)	Mean share CNP
Quartile	Mean				
Q1	0.129	0.000	76.2	2,722.1	0.402
Q2	0.320	0.001	66.6	3,012.2	0.476
Q3	0.492	0.004	54.2	3,332.3	0.513
Q4	0.672	0.063	45.5	3,056.4	0.416

Note: Table shows descriptive statistics on accounts grouped by quartiles of the predicted probability that they used cash advance. The probability of using cash advance is predicted using a random forest classifier, described in Appendix B. For each quartile, we show the mean predicted probability of having a cash advance, the actual share of accounts with a cash advance, the mean number of monthly card transactions, the mean monthly card spending, and the mean share of card-not-present (CNP) transactions.